AFRL-RQ-WP-TR-2019-0215

GEOMETRIC OPTIMIZATION AND COMBINATORIAL-HOMOLOGICAL PROGRAMMING

Robert Ghrist
Trustees of The University of Pennsylvania
The Clinical Practices of The University of Pennsylvania

OCTOBER 2019
Final Report

DISTRIBUTION STATEMENT A. Approved for public release.
Distribution is unlimited.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
AEROSPACE SYSTEMS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE

NOTICE AND SIGNATURE PAGE

AFRL-RQ-WP-TR-2019-0215 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

BERAN.PHILIP.S.1020844490 Digitally signed by BERAN.PHILIP.S.1020844490 Date: 2020.01.27 10:12:18 -05'00'

PHILIP S. BERAN, PhD
Work Unit Manager and Technical Advisor
Design and Analysis Branch

TYLER.CHARLES.1229064106 Digitally signed by TYLER.CHARLES.1229064106 Date: 2020.01.28 08:39:23 -05'00'

CHARLES TYLER
Chief, Design and Analysis Branch
Aerospace Vehicles Division

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
October 2019	Final	3 April 2018 – 3 October 2019

4. TITLE AND SUBTITLE
GEOMETRIC OPTIMIZATION AND COMBINATORIAL-HOMOLOGICAL PROGRAMMING

5a. CONTRACT NUMBER
FA8650-18-2-7840

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
61101E

6. AUTHOR(S)
Robert Ghrist

5d. PROJECT NUMBER
CCS-02 (DARPA)

5e. TASK NUMBER

5f. WORK UNIT NUMBER
Q1Y9

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Trustees of The University of Pennsylvania
The Clinical Practices of The University of Pennsylvania
3451 Walnut Street
Philadelphia, PA 19104-6205

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Aerospace Systems Directorate
Wright-Patterson Air Force Base, OH 45433-7542
Air Force Materiel Command
United States Air Force

Defense Advanced Research Projects Agency/Defense Sciences Office (DARPA/DSO)
3701 N. Fairfax Drive
Arlington, VA 22203

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/RQVC

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-RQ-WP-TR-2019-0215

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

13. SUPPLEMENTARY NOTES
This report is the result of contracted fundamental research, which is deemed exempt from Public Affairs Office security and policy review in accordance with Deputy Assistant Secretary of the Air Force (Science, Technology, Engineering) (SAF/AQR) memorandum dated 10 Dec 08 and Air Force Research Laboratory Executive Director (AFRL/CA) policy clarification memorandum dated 16 Jan 09.

14. ABSTRACT
New mathematical methods for distributed optimization and related problems were generated using concepts from algebraic topology and sheaf theory. In particular, the Hodge Laplacian was extended to sheaves of data in the context of distributed optimization, indoor mapping, network communications, and learning. The results greatly extend known results from spectral graph theory and simultaneous localization and mapping.

15. SUBJECT TERMS
optimization, topology, spectral graph theory, sheaves, Laplacian

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT:	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON (Monitor)
Unclassified	Unclassified	Unclassified	SAR	21	Philip S. Beran
					19b. TELEPHONE NUMBER (Include Area Code) N/A

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18

Table of Contents

List of Figures

1. Summary

Linear, convex optimization is an economically and strategically critical industry whose efficient solutions impact every aspect of modern transportation, commerce, communication, and prediction. Current demands – from artificial intelligence (AI) to robotics, autonomous vehicles, networks, neuroscience, and more – prompt increasingly high-dimensional, non-linear, and non-convex problems with an urgency that outstrips our present growth rate in computational abilities.

1.1 Three Themes

The project involved development of new methods in optimization theory along three components of a unified whole:

1) **Homological Programming :** injecting homological algebra as a tool to constrain or collate optimization problems;
2) **Spectral Sheaf Theory :** lifting spectral graph theory from networks to sheaves and balancing geometric, homological, and spectral data of sheaf Laplacians; and
3) **Cubification and Curvature :** adapting geometric group theory and CAT(0) cubical complexes to quickly solve optimization problems on complex nonconvex domains.

These represent distinct aspects of a unified revolution in optimization based on homological-algebraic methods in geometric, combinatorial, and analytic emanations.

1.2 Goals

This project consisted of fundamental research. The goal of this project was to produce a revolutionary set of tools and perspectives with which to address optimization problems that are potentially high-dimensional, distributed, and/or involving nonstandard data types.

The approach of this project was not in the direction of a generalization of existing methods; rather, the work here summarized pushed out into new and revolutionary methods that have previously had little to no impact in optimization theory (such as homological algebra, geometric group theory, and sheaf theory).

2. Homological Programming

2.1 Distributed Optimization and Homological Constraints

Our goal of defining homological programs as a globally constrained set of distributed optimization problems with homological constraints evolved to the following general framework. Consider an undirected graph G with sheaf $\mathcal{H}$ of vector spaces over vertices and edges of G. Consider a set of convex optimization problems indexed by vertices v of G; that is, $F = \{f_v\}$, where each f_v is a convex functional on the stalk $\mathcal{H}_v$.

Problem: ([HG19-1] §II) *Minimize $F(x) = \sum_v f_v(x_v)$ subject to the global constraint that x is a global section of the sheaf $\mathcal{H}$.*

We used a primal-dual method with the sheaf Laplacian $L_{\mathcal{H}}$ to descend on the primal variable x and ascend on the dual variable z. This led to the following general result:

Theorem: (*cf.* [HG19-1] §IV) *Flowing according to the sheaf Laplacian via the dynamical system*

$$\frac{dx}{dt} = -\nabla f(x) - 2L_{\mathcal{H}}(x) - L_{\mathcal{H}}(z) \quad : \quad \frac{dz}{dt} = L_{\mathcal{H}}(x)$$

converges to the Karush-Kuhn-Tucker [KKT] points of the optimization problem above.

We generated several applications of this new technique: see [HG19-1] §V, which established the theory and outlined applications to:

1. Distributed optimal control of linear systems
2. Distributed optimal signal reconstruction on graphs

2.2 Application: Linear Optimal Controls

One application to a classical setting is in linear control systems. Consider the following standard time-dependent linear system with controls on $\mathbb{R}^n$:

$$x_{t+1} = A_t x_t + B_t u_t \; ; \; y_t = C_t x_t + D_t u_t$$

This can be represented as a sheaf over a tree as in Figure 1.

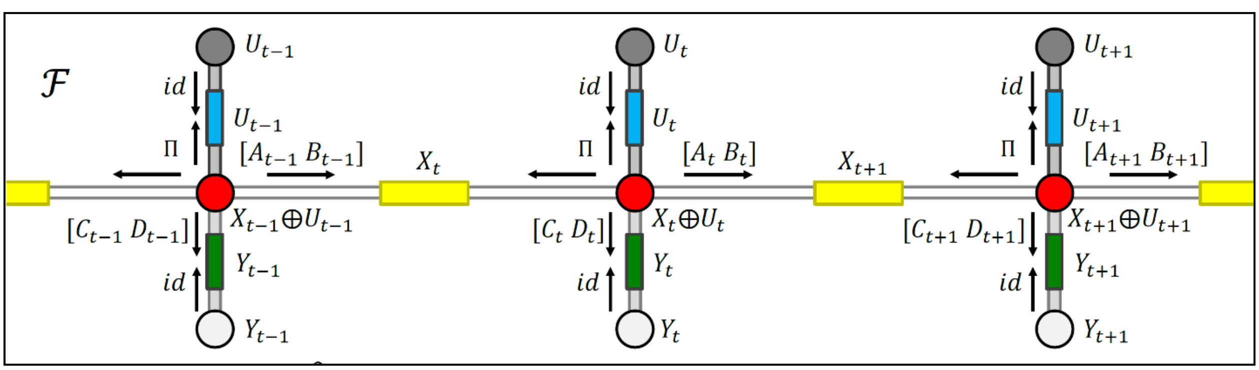

Figure 1: A linear control system is realized as a sheaf of vector spaces over a tree.

Here, id means the identity map and Π means projection onto the appropriate factor. The sheaf Laplacian $L_{\mathcal{F}}$ then converges to $H^0(\mathcal{F})$, the global-in-time solutions to the system. Fitting this into the framework of homological programming, we can implement sheaf Laplacian evolution to find global solutions to the linear system which optimize (time-dependent) functions.

2.3 Scenario: Indoor Mapping through Nonco-Operative Sensing

As a testbed for our qualitative methods, we have built the following scenario. Consider a setting in which control has been exerted over multiple security camera feeds in a building.

2.3.1 Assumptions

The following capabilities or sensor inputs are assumed:

1. Target ID, locally consistent in time. Persons in view can be tracked while in-view and identities can be re-established when targets appear on another feed.
2. Consistent timing between feeds, so that it is known how much time has passed between targets entering and exiting field-of-view.
3. Local clustering of target entrance/exit regions within a camera's field of view.

Things that are ***not*** assumed include building layout data, location or bearing of cameras, distribution or density of cameras, or any other geographic data.

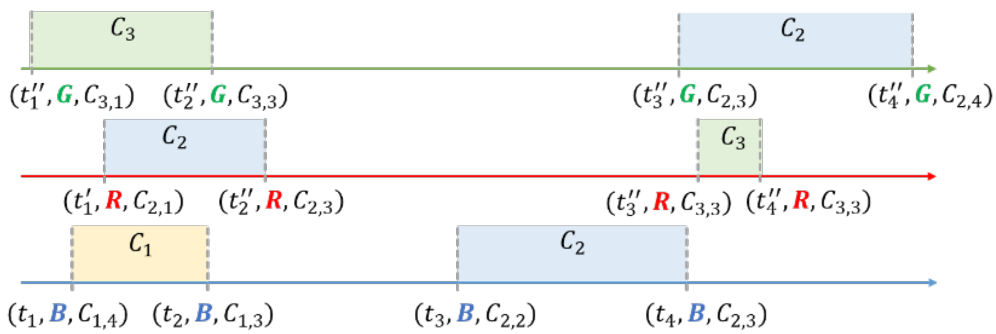

Figure 2: Agents are tracked in time according to visibility in-field.

From this data, the goal is to build: 1) a topologically-accurate map of the interior layout of the building, getting details like hallway configuration correct up to homology; 2) a geometric relaxation based on average transit times, as an approximation to the domain geometry.

2.3.2 Technique

Our technique involved the following ingredients:

1. Clustering entrance/exit domains in camera fields-of-view. These clusters will form the vertex set V of a combinatorial model of the domain topology.
2. Use of average minimal transit times between elements in V to generate a weighted transit matrix T.
3. Building a modified ***Dowker complex*** of T as a simplicial cell complex model of the domain, with guarantees on homological accuracy.
4. Augmenting the model with refinements of the transit time distributions in order to determine "hidden hallways" that no camera has knowledge of.
5. Setting up an SDP whose solution gives the topological model a geometric approximant.

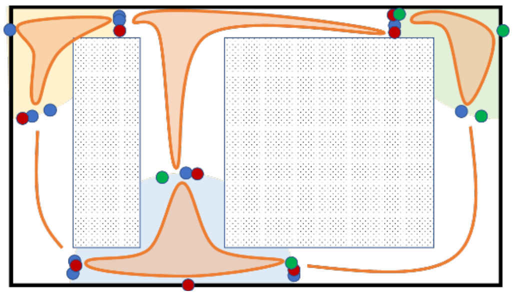

Figure 3: Topological reconstruction of floorplan is obtained.

This can be seen in the following simple example, where a simple floor plan is surveilled with three cameras. Agents moving in and out of the field of view are tracked on a time line [Figure 2]. Clustering entrance and exit regions and building the Dowker complex of the transit matrix yields a simplicial cell complex which captures the correct homotopy type of the floorplan [Figure 3].

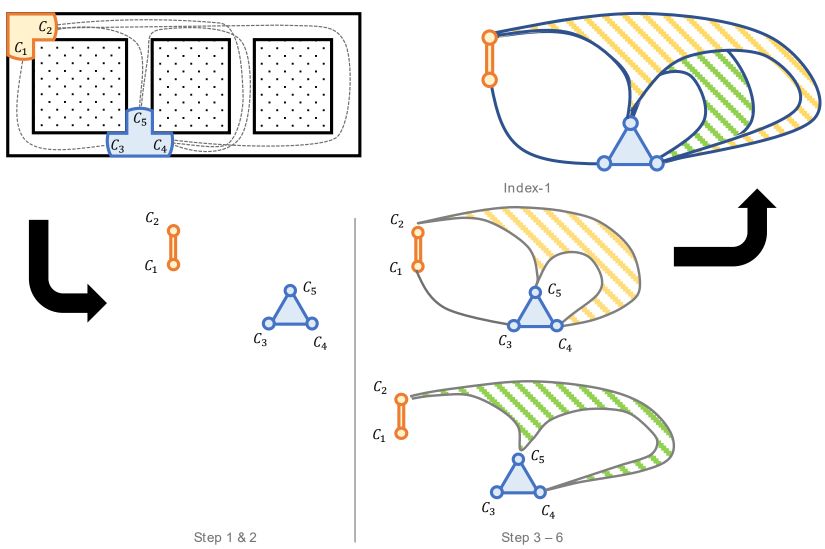

Figure 4: Floorplan reconstruction is possible from sparsely-placed cameras.

Even in situations where the floorplan is fairly complicated and the systems of sensors is sparse, one can still provably reconstruct the proper topology of the floorplan. Figure 4 gives an example of two cameras which suffice to reconstruct a complex floorplan. The key is to control the contractability of the complement of the covered region.

With the ability to approximate indoor map topology and geometry of a single floor, we turned to the multi-floor setting. Assume the following constraints: floors are connected by elevators (perhaps in various locations) and the elevators are *covered* – there are cameras (whether inside the elevators or at entrance/exit doors) which indicate who gets in/out when.

The problem is to infer 1) the number of floors; 2) the total ordering of the floors; and 3) an identification of which cameras are observing which floors, in order to reconstruct the entire building floorplan topology.

With assumptions as stated, along with a “efficient rider” assumption (people take the shortest elevator path when changing floors and, e.g., do not go in an up elevator when they meant to go down), these problems are solvable.

- Step 1 is to use the elevator coverage assumption to partition all remaining cameras into connected floors (with unknown floor numbers).
- Step 2 is to use accumulated observations about entrances/exits from elevators along with the no-idiots assumption to determine the terminal floors (top/bottom). Top can be distinguished from bottom based on circadian patterns (when most people arrive or depart for the day).
- Step 3 is to use interleaving of elevator entrances/exits with Step 2 to obtain the ordering of the floors.

2.3.3 Simulation Engine

In order to test the effectiveness of the methods above, we built a simulation engine [Figures 5 and 6]. Using a combination of python programming within the popular *Unity* game engine environment, we set up a procedurally-generated building with customizable agents wandering through it. Besides procedurally-generated floorplans, this simulation engine has a customizable set of cameras and placement parameters, adjustable speed and behaviors of agents, and variable time-scales (for quick-run data collection).

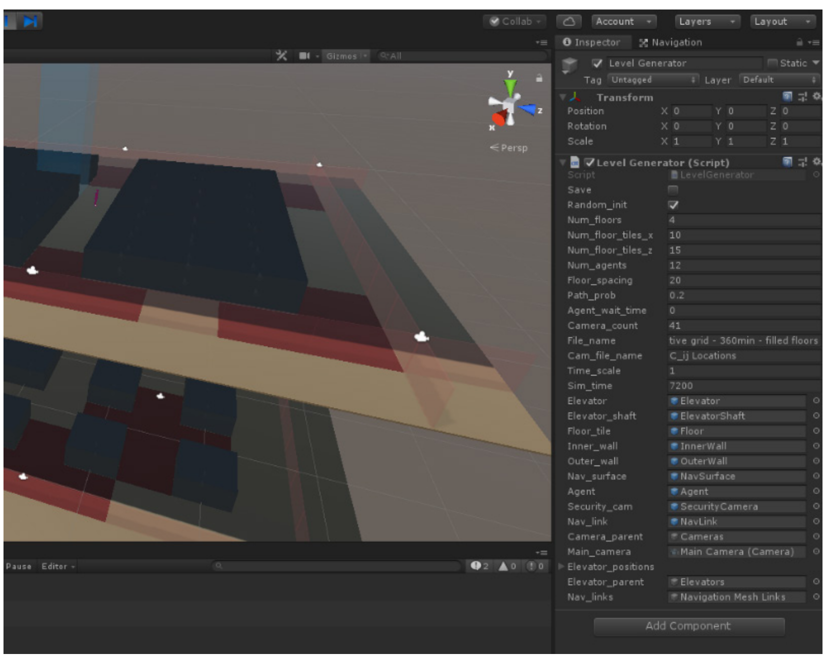

Figure 5: Multi-floor buildings are simulated.

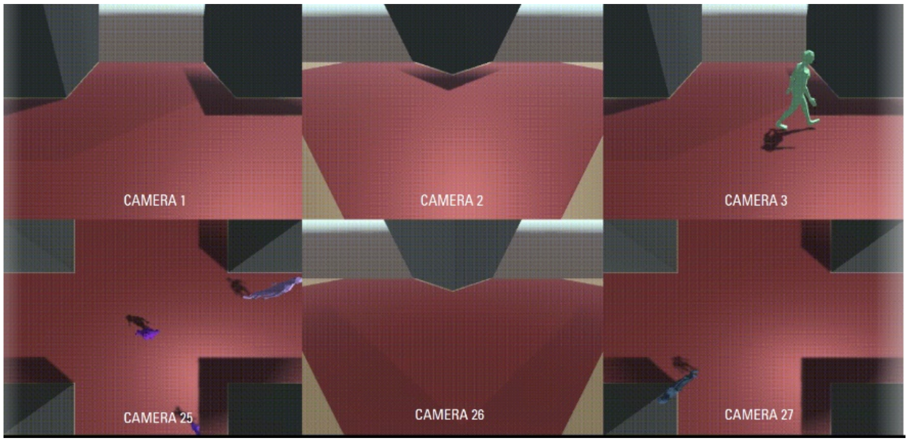

Figure 6: Still shots are shown from the monitored cameras in the simulation engine.

2.3.4 Analysis and Summary

Proposition III.3 of [GS19] proves that the process described above produces a 1-d cell complex which, in the case of a planar domain with unimodal transit times between cameras with simply-connected complements, returns the correct topology of the floorplan. Results in Sections IV and V of [GS19] detail the use of SDPs to approximate geometric information. No hard results are provable due to the weakness of the assumptions. As there are no other methods that work with the minimal-sensing model we have employed, no comparisons to existing techniques can be made, other than to note that the topological techniques return provably correct results with virtually no assumptions on the quality of sensor input. Future work should focus on the geometric approximation: [GS19] indicated that adding constraints to the SDPs (such as the existence of known anchor points) improved results: quantifying the degree of improvement possible would be an avenue for future results.

2.4 Learning Sheaves

The sheaf Laplacian can be used to define a notion of consistency – or *smoothness* – of data on the vertices. In the same way that one can learn graph structures from signals over the vertices which are smooth with respect to the graph Laplacian, we showed how to learn a sheaf based on vertex-signals smooth with respect to the sheaf Laplacian [Figure 7].

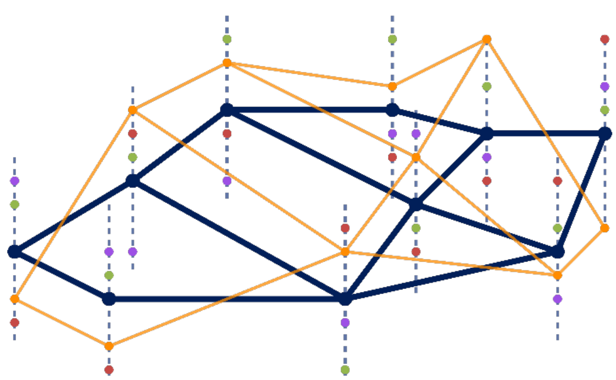

Figure 7: Multimodal signals over a graph can be learned.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This is significant in that the signals can be multimodal and quite different from vertex to vertex. The solution technique here, as with the classical case of graph Laplacians, is to frame the problem as a constrained optimization problem. The unique feature here is the global constraint that makes this an excellent example of homological programming. This work is written up in [HG19-2].

2.5 Sheaves of Opinions and Opinion Dynamics on Networks

A class of ***opinion sheaves*** meant to model how individuals negotiate ideological agreement and disagreement on a variety of issues that are context-dependent and mediated through a social network was initiated. Stalks $\mathcal{F}_v$ are represented as positive cones in a real vector space, whose basis vectors represent, roughly, the set of issues on which agent v has a strongly-held opinion on a subject (public policy, preference, politician, etc.) with binary outcome (like/dislike) coupled with degree. Stalks over edges $\mathcal{F}_e$ are (large) vector spaces spanned by all possible issues/preference questions. The restriction maps for the sheaf are maps of positive cones into the net opinion vector space. These maps perform localized opinion alignment.

One important feature of these sheaves are that the maps $\mathcal{F}_v \to \mathcal{F}_e$ are not required to be identity maps (composed with the standard embedding). For example, the restriction maps can reverse orientation along certain axes (this is, a person can lie about their opinion to certain others). This increases greatly the ability to model complex social interactions.

It was shown that a diffusion operator built from the sheaf Laplacian suffices to bring such a system into global consensus, if such a consensus exists: more formally, the fixed point set of the diffusion operator is isomorphic to the space of global sections of the sheaf. A preliminary report appears in [HG18] §8.

2.6 Sheaves of Lattices and Laplacians

Work was done on the use of dynamics to compute global solutions to local systems of constraints. At first, this involved using dynamics on *lattices of cochain complexes*, but was generalized to the similar (but vastly broader) setting of *cochain complexes of lattice*s. A *lattice* is a partially-ordered set (or *poset*) outfitted with algebraic binary operations of *meet* (written $\wedge$, acting like an intersection or greatest lower bound) and *join* (written $\vee$, acting like a union or least upper bound). Specific examples include Boolean algebras and similar structures from logic.

The goal was to prove convergence results that a given dynamical iterative scheme can be proven to converge to the cohomology of the underlying chain complex. This is related to the project on spectral sheaves, where the Laplacian and a global heat equation provides the canonical example of such.

A class of cellular sheaves $\mathcal{F}$ of lattices over graphs was considered: data structures over a network where to each vertex and edge is assigned an abstract lattice, and to each vertex-edge pair (v, e) is assigned a lattice morphism $\mathcal{F}_{v \le e}: \mathcal{F}_v \to \mathcal{F}_e$ between the corresponding lattices. Using the Galois correspondence leads to an adjoint lattice morphism $\mathcal{F}^*_{v \le e}: \mathcal{F}_e \to \mathcal{F}_v$.

A novel Laplacian operator acting on the union of the vertex-data lattices was discovered. This is defined on vertex stalks in [RG19] as follows:

$$\Delta(x)_v := \bigwedge_{v-e-w} \mathcal{F}^*_{v\leq e}\big(\mathcal{F}_{v\leq e}(x_v) \wedge \mathcal{F}_{w\leq e}(x_w)\big)$$

The intuition for why this is a Laplacian stems from the flow of data up from vertices to edges, then back down to vertices, much like a graph Laplacian.

A diffusion process using this Laplacian in a discrete-time flow was developed, much as what occurs when solving a heat equation on a domain, to relax an initial condition to a convergent solution. Using this operator, the following result was obtained: a perfect analogue of the discrete-time Laplacian flow converging to harmonic solutions.

Theorem: ([RG19]) *For $\mathcal{F}$ a sheaf of (complete meet semi-)lattices of a graph G, then the discrete time Laplacian evolution operator*

$$\Phi := \Delta \wedge Id : C^0(G;F) \to C^0(G;F)$$

Is an endomorphism on 0-cochains with fixed point set $Fix(\Phi) = H^0(G;\mathcal{F})$ the space of global sections of the sheaf. Moreover, if the stalks of $\mathcal{F}$ are finite, then Φ converges to global sections in a finite number of iterations.

This has as immediate corollary the ability to do distributed consensus on networks of fixed lattices. This in itself is interesting, but is not nearly as powerful as the ability to evolve over a sheaf of lattices. This work is currently being extended and written up for publication [RG19].

2.7 Application: Wireless Communication Networks

An interesting application of sheaves of lattices was found for a classical problem in wireless communications. The typical (conservative) protocol for wireless network communications follows these steps:

1. RTS : Request To Send : node A broadcasts "Request to talk to B!"
2. CTS : Clear To Send : when everyone goes silent, B says "OK!"
3. DAT : Data transfer : A broadcasts data read by B
4. ACK : Acknowledgement : B says "Message received!"

This is clearly over-conservative and redundant.

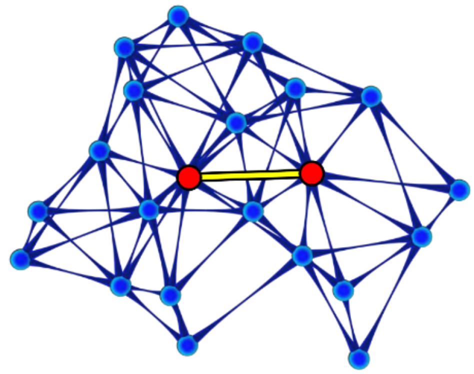

Figure 8: The 2-hop neighborhood of an edge in a local network is illustrated.

We constructed a sheaf of (meet semi-)lattices over the communication network that encodes legal communication protocols. For each edge e in the graph, consider the 2-hop neighborhood of edges incident to e [Figure 8]. The stalk over e is the semi-lattice of subsets of edges in the 2-hop neighborhood generated by sets of the form 1) any subset not containing e (representing cases where e is "off"; and 2) subsets containing e for which communication over said edges does not overwhelm communication over e. Stalks over vertices are the full lattice of subsets of edges incident. The restriction maps glue vertex data into edge data via embeddings.

Proposition: [RG19] *The zeroth cohomology H^0 of this sheaf consists of all legal communications scenarios in which signal-interference does not impede communication.*
Much more remains to be done on this problem, including determining how to project from an initial, non-sectional state of communication, to a collection of legal local sections for the sheaf in order to apply a Laplacian smoothing to an optimal communications arrangement for continuous communication without interruption.

3. Spectral Sheaf Theory

3.1 Sheaf Laplacians and Spectra

This project has led to the development of a vast generalization of spectral graph theory to cellular sheaves – data structures over cell complexes taking values in vector spaces. To such, one can associate Laplacians (of various sorts), the spectra of which carry information about the sheaf as well as the underlying complex. In spectral graph theory, one is really examining the spectra of the up-Laplacian of the constant sheaf over a graph. This is a huge area in applied mathematics, with applications to optimization, clustering, graph sparsification, and much more. The work done during this project greatly generalizes this subject and has opened multiple new directions for theoretical and applied investigations.

Our inquiries led to a 35-page paper [HG18] which contained the following:

1. Definitions of Laplacians (standard, up- and down-) associated to sheaves of vector spaces over cell complexes.
2. Characterization of spectra of cellular sheaves under the operations of pushforwards of covering maps and sheaf products.
3. Examples to demonstrate multiple non-isomorphic sheaves with the same Laplacians: i.e., one cannot "hear" the shape of a sheaf.
4. Examples of obstructions to harmonic extension (cf. Kron reduction) on sheaves.
5. Definitions of effective resistance for cycles on cosheaves over cell complexes.
6. A theorem on computability of effective resistance in terms of the pseudo-inverse of the up-Laplacian.
7. A sparsification theorem for cellular sheaves.
8. A roadmap for spectral sheaf theory applications, including synchronization, consistent clustering of data, sheaf approximations, and more.

3.2 Effective Resistance and Sampling

The definition of effective resistance is of importance: the idea is that the effective resistance of a pair of homologous cycles is the minimal cobounding chain that spans them. So, for (α, β) a pair of k-dimensional cycles in a complex,

$$R_{eff}(\alpha,\beta) = \min_{\gamma} |\gamma| \ : \ \gamma \in C^{k+1}(X;\mathcal{F}) \ : \partial\gamma = \beta - \alpha$$

This is shown to be computable in terms of the pseudo-inverse of the up-Laplacian.

The most significant theorem to come about is the sparsification theorem for sheaves, generalizing a well-known result in spectral graph theory due to Spielman et al. [ST11].

***Theorem*:** ([HG18]) *Let X be a regular cell complex of dimension d and $\mathcal{F}$ a cosheaf on X with $dim\, C_{d-1}(X;\mathcal{F}) = n$. Given $\epsilon > 0$ there exists a subcomplex $X' \subset X$ with the same $(d-1)$-skeleton and $O(\epsilon^{-2} n \log n)$ d-cells, together with a cosheaf $\mathcal{F}'$ on X' such that the Laplacians are ϵ-close:*

$$(1-\epsilon)L_{\mathcal{F}}^{d-1} \preccurlyeq L_{\mathcal{F}'}^{d-1} \preccurlyeq (1+\epsilon)L_{\mathcal{F}}^{d-1}$$

The proof of this works via probabilistic sampling according to effective resistance.

3.3 Sheaf Approximation Theorems

One of the issues that has arisen is the need to approximate a (high-dimensional, high-rank) sheaf with a simpler model that captures the "essential" information in the original sheaf. The natural notion in light of this program is that of a *spectral* approximation: the approximating sheaf has smaller-dimensional stalks but has nearby nonzero spectrum in its Laplacian. This leads naturally to questions of expanders.

3.4 Expander Sheaves

One of the more important tools in applications of spectral graph theory to error-correcting codes, embedding problems, and even Riemannian geometry comes from the notion of *expander graphs* [HLW06]. An expander in a graph is a subgraph which "feels hyperbolic" in the sense of being well-connected with large boundary relative to interior.

Motivated by a conjectural Cheeger inequality for sheaves over graphs, we extended the notion of spectral expander graphs to certain classes of sheaves over graphs, namely *quasi-constant* sheaves, in which all vertex stalks are a fixed vector space and restriction maps are constant on each edge (but can vary from edge to edge). Among such quasi-constant sheaves, an η*-expander* is k-regular (the Laplacian is block-diagonal with entries kI) and whose adjacency matrix has spectrum bounded by η.

Our principal result here is an expansion of the well-known *Expander Mixing Lemma* for expander graphs [HLW06], which says that for an η-expander graph G, the expected number of edges between subsets $S, T \subset V(G)$ differs from the known k-regular mean by

$$\left|\mathbb{E}\big(E(S,T)\big) - \frac{k|S||T|}{|G|}\right| \leq \eta\sqrt{|S||T|\left(1-\frac{|S|}{|G|}\right)\left(1-\frac{|T|}{|G|}\right)}.$$

To extend this result to sheaves, we have shown the following:

Lemma: [H20] *For a k-regular quasi-constant sheaf over a graph G whose vertex stalks have dimension d, and for subsets $S, T \subset V(G)$,*

$$\left|tr\big(E(S,T)\big) - \frac{kd|S||T|}{|G|}\right| \leq \sum_{i=1}^{d}|\lambda_{d+i}|\sqrt{|S||T|\left(1-\frac{|S|}{|G|}\right)\left(1-\frac{|T|}{|G|}\right)},$$

where tr represents the traces of the appropriate blocks and the sheaf adjacency matrix has ordered eigenvalues $\lambda_1 = \cdots = \lambda_d \leq \lambda_{d+1} \leq \cdots$

Connections between this work and the isoperimetric inequality results incident to a Cheeger inequality have been established.

3.5 Marginalization Sheaves

As a potential application for the spectral theory of the sheaf Laplacian, a new class of cellular sheaves over a network was generated. Motivated by graphical models in machine learning and inference, we have developed the following class of *marginalization sheaves*.

Given a collection $\{X_i\}$ of discrete random variables with factor graph G, one knows that the joint pdf factors according to G. The problem is, given the potential functions of the individual factors, compute the local marginals (without knowing the joint pdf, since there are potentially so many variables involved) [WJ08].

The marginalization sheaf is a sheaf over the Hasse diagram H of the poset derived from the factor graph G. The stalk over a vertex/edge is the vector space of all possible probability distributions over the implicated variables, and the restriction maps of the sheaf are marginalizations. Results proved during the program imply that in the case where the base network H is acyclic, then Laplacian evolution realizes a unique solution to the problem as an optimization involving Shannon entropy. At this time, it is not clear whether this method is essentially the same as the traditional message-passing techniques used to optimize in a distributed manner.

4. Cubification and Curvature

4.1 Cubings

Numerous problems in optimization (such as the shortest path problem) are sensitive to the curse of dimensionality: shortest path goes from complexity P to complexity NP-hard as one goes from polygonal domains in 2-d to 3-d. One method of attack is to restrict the geometry of the underlying domain. For example, in shortest path, restricting from arbitrary domains to ***CAT(0)*** cube complexes reduces NP-hard back to P; a truly hyperbolic geometry further reduces complexity to linear.

Our goal was to simultaneously learn a space while optimizing a function on it. The only hope for doing so efficiently is to similarly restrict to the setting of nonpositively curved (***NPC***) or locally CAT(0) cubings – spaces built from cubes arranged in such a way as to eliminate any positive curvature. Such spaces have been shown to be common in robotics, planning, phylogenetics, and other reconfigurable information systems.

4.2 Universal Memory Architectures

Our work was related to an in-progress architecture for learning the NPC cubing. Such ***universal memory architectures*** (UMAs) are highly efficient with regards to memory allocation and can be used to learn the lattice (more precisely ***pocset***) structure underlying an NPC cubing. We have established the preparatory results on UMAs (how they work, what they can do, comparison to SOA methods such as neural networks). Current results on optimization include:

1. For an integer-valued function on the vertex set of an NPC cubing, a UMA can efficiently learn the convex hull of the optima.
2. We have a set of conditions on convexity of intersections in a cosheaf of CAT(0) cubings that allows us to adapt UMAs for learning.

This prerequisite background work on UMAs was written up in a 58-page manuscript [GK18]. We proceeded to cast the problem of cubical complex reconstruction from distributed sensor data as an optimization problem using generalized Laplacians. Specifically, let X be a CAT(0) cubical complex and G its 1-skeleton graph. The structure of X is unknown: only knowledge of G is available, and that only through a time-series of observations as follows.

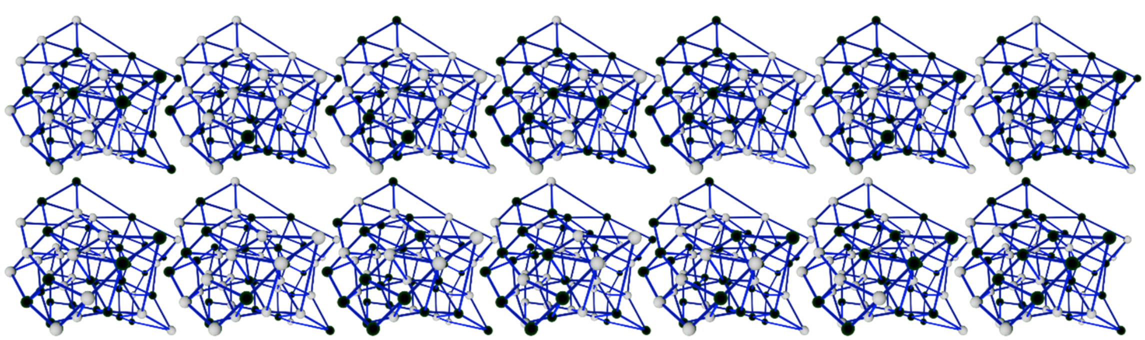

Figure 9: Shown is a split system of states on vertices of an unknown cubical complex.

Consider a binary functional s from the vertex set of the graph, $V(G)$, to the set $\{0,1\}$. A system $S = \{s\}$ of such binary functionals is said to be a *split system* if each s is non-constant and if S is closed under complements and distinguishes vertices.

Conjecture: *With S a split system as above, the functional*

$$\Psi(S) = \sum_{s \in S} |\delta(s)| - \frac{1}{2}|S|$$

acts as a weighted Laplacian whose minimizers yield the hyperplane structure for the cubical complex X.

A partial result was proved during this program: it is a minimax theorem of the following form:

Proposition: [GG19] *Maximizers of $|S|$ subject to the constraint that $\sum_{s \in S} |\delta(s)|$ is minimized yields the hyperplane structure of the underlying cubical complex.*

5. Summary and Conclusions

Prior to this project, the applications of algebraic topology to optimization theory were virtually nonexistent. This program developed several novel methods for using algebraic topology in distributed optimization, among which included:

- The first applications of the Hodge Laplacian of cellular sheaves.
- A revolutionary extension of spectral graph theory to a spectral theory of cellular sheaves and cosheaves.
- Definition of Homological Programming in terms of optimization with cohomological constraints.
- Creation and analysis of Laplacian diffusion techniques for homological programs, with convergence results.
- Extension of Laplacian methods to sheaves taking values in certain categories of posets and lattices.
- Adoption of Laplacian optimization methods to Universal Memory Architectures and reconstruction of cubical configuration spaces from sampled time-series data.

In addition, several applications were outlined: including to problems of control, clustering, network communication, indoor mapping, learning, and opinion dynamics on social networks. Though the focus of the project was on development of fundamental theory, several of the initial applications appear promising enough to warrant future attention. In particular, the indoor-mapping problem, which utilized minimal uncooperative sensing from a set of cameras tracking agents, yielded topologically-inferred maps unavailable by any other known means from the minimal input data.

6. References

[GG19] D. Guralnik and R. Ghrist, "Identifying a finite cubing from Boolean queries", in preparation.

[GK18] D. Guralnik and D. Koditschek, "Iterated Belief Revision Under Resource Constraints: Logic as Geometry", ArXiv preprint arXiv:1812.08313

[H20] J. Hansen, "Expander Sheaves", in preparation.

[HG18] J. Hansen and R. Ghrist, "Toward a Spectral Theory of Cellular Sheaves", *Journal of Applied & Computational Topology*, DOI:10.1007/s41468-019-00038-7, 2019.

[HG19-1] J. Hansen and R. Ghrist. Distributed Optimization with Sheaf Homological Constraints. *Allerton Conference on Communication, Control, and Computing*, 2019.

[HG19-2] J. Hansen and R. Ghrist, "Learning sheaf Laplacians from smooth signals", *IEEE Conf. on Acoustics, Speech, & Signal Processing [ICASSP]*, to appear.

[GS19] R. Ghrist and A. Speranzon, "Topological mapping with uncooperative sensing", to appear, Proc. ICRA.

[HLW06] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications", in *Bull. Amer. Math. Soc.,* 43(4): 439-561, 2006.

[RG19] H. Riess and R. Ghrist, "Laplacian Diffusion on Sheaves of Lattices", in preparation.

[ST11] D. Spielman and S.-H. Teng. "Spectral Sparsification of Graphs," *SIAM Journal on Computing*, 40(4):981–1025, 2011.

[WJ08] M. Wainwright and M. Jordan. "Graphical Models, Exponential Families, and Variational Inference", in *Found. Trends Mach. Learn.*, 1(1-2):1–305, 2008.